IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
Before the Board of Patent Appeals and Interferences

In re Patent Application of

HE

Serial No. 10/507,112

Filed: September 10, 2004

Atty Dkt. CC-36-1842
C# M#
Confirmation No. 4135
TC/A.U.: 2129
Examiner: O. Fernandez Rivas
Date: August 4, 2008

OIPE IAP38 AUG 0 4 2008 PATENT & TRADEMARK OFFICE

Title: METHOD OF COMBINATORIAL MULTIMODAL OPTIMISATION

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

☐ **Correspondence Address Indication Form Attached.**

☐ **NOTICE OF APPEAL**
Applicant hereby **appeals** to the Board of Patent Appeals and Interferences from the last decision of the Examiner twice/finally rejecting applicant's claim(s). $510.00 (1401)/$255.00 (2401) $

☒ An appeal **BRIEF** is attached in the pending appeal of the above-identified application $510.00 (1402)/$255.00 (2402) $ 510.00

☐ Credit for fees paid in prior appeal without decision on merits -$ ( )

☐ A reply brief is attached. **(no fee)**

☐ Petition is hereby made to extend the current due date so as to cover the filing date of this paper and attachment(s)
One Month Extension $120.00 (1251)/$60.00 (2251)
Two Month Extensions $460.00 (1252)/$230.00 (2252)
Three Month Extensions $1050.00 (1253/$525.00 (2253)
Four Month Extensions $1640.00 (1254/$820.00 (2254) $

☐ "Small entity" statement attached.

Less month extension previously paid on -$( )

TOTAL FEE ENCLOSED $ 510.00

☐ **CREDIT CARD PAYMENT FORM ATTACHED.**

Any future submission requiring an extension of time is hereby stated to include a petition for such time extension. The Commissioner is hereby authorized to charge any deficiency, or credit any overpayment, in the fee(s) filed, or asserted to be filed, or which should have been filed herewith (or with any paper hereafter filed in this application by this firm) to our **Account No. 14-1140.** A duplicate copy of this sheet is attached.

901 North Glebe Road, 11th Floor
Arlington, Virginia 22203-1808
Telephone: (703) 816-4000
Facsimile: (703) 816-4100
CC:lmr

NIXON & VANDERHYE P.C.
By Atty: Chris Comuntzis, Reg. No. 31,097

Signature: ______________

1366883

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
Before the Board of Patent Appeals and Interferences

In re Patent Application of

HE

Serial No. 10/507,112

Filed: September 10, 2004

Atty Dkt. CC-36-1842
C# M#
Confirmation No. 4135
TC/A.U.: 2129
Examiner: O. Fernandez Rivas
Date: August 4, 2008

OIPE IAP38 AUG 0 4 2008 PATENT & TRADEMARK OFFICE

Title: METHOD OF COMBINATORIAL MULTIMODAL OPTIMISATION

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

☐ **Correspondence Address Indication Form Attached.**

☐ **NOTICE OF APPEAL**
Applicant hereby **appeals** to the Board of Patent Appeals and Interferences from the last decision of the Examiner twice/finally rejecting applicant's claim(s). $510.00 (1401)/$255.00 (2401) $

☒ An appeal **BRIEF** is attached in the pending appeal of the above-identified application $510.00 (1402)/$255.00 (2402) $ 510.00

☐ Credit for fees paid in prior appeal without decision on merits -$ ( )

☐ A reply brief is attached. **(no fee)**

☐ Petition is hereby made to extend the current due date so as to cover the filing date of this paper and attachment(s)
One Month Extension $120.00 (1251)/$60.00 (2251)
Two Month Extensions $460.00 (1252)/$230.00 (2252)
Three Month Extensions $1050.00 (1253/$525.00 (2253)
Four Month Extensions $1640.00 (1254/$820.00 (2254) $

☐ "Small entity" statement attached.

Less month extension previously paid on -$( )

TOTAL FEE ENCLOSED $ 510.00

☐ **CREDIT CARD PAYMENT FORM ATTACHED.**

Any future submission requiring an extension of time is hereby stated to include a petition for such time extension. The Commissioner is hereby authorized to charge any deficiency, or credit any overpayment, in the fee(s) filed, or asserted to be filed, or which should have been filed herewith (or with any paper hereafter filed in this application by this firm) to our **Account No. 14-1140.** A duplicate copy of this sheet is attached.

901 North Glebe Road, 11th Floor
Arlington, Virginia 22203-1808
Telephone: (703) 816-4000
Facsimile: (703) 816-4100
CC:lmr

NIXON & VANDERHYE P.C.
By Atty: Chris Comuntzis, Reg. No. 31,097

Signature:

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

OIPE IAP38
AUG. 0 4 2008
PATENT & TRADEMARK OFFICE

In re Patent Application of

HE — Atty. Ref.: 36-1842

Serial No. 10/507,112 — TC/A.U.: 2129

Filed: September 10, 2004 — Examiner: O. Fernandez Rivas

For: METHOD OF COMBINATORIAL MULTIMODAL OPTIMISATION

* *

August 4, 2008

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

APPEAL BRIEF

Sir:

Appellant hereby **appeals** to the Board of Patent Appeals and Interferences from the last decision of the Examiner.

08/05/2008 JADDO1 00000066 10507112
01 FC:1402 510.00 OP

TABLE OF CONTENTS

(I) REAL PARTY IN INTEREST

The real party in interest is British Telecommunications public limited company, a corporation of the country of the United Kingdom.

(II) RELATED APPEALS AND INTERFERENCES

The appellant, the undersigned, and the assignee are not aware of any related appeals, interferences, or judicial proceedings (past or present), which will directly affect or be directly affected by or have a bearing on the Board's decision in this appeal.

(III) STATUS OF CLAIMS

Claims 1, 3-21, and 23-25 are pending and have been rejected. Claims 2 and 22 have been cancelled. No claims have been substantively allowed. All of rejected claims 1, 3-21, and 23-25 are being appealed.

(IV) STATUS OF AMENDMENTS

No amendments have been filed since the date of the Final Rejection.

(V) SUMMARY OF CLAIMED SUBJECT MATTER

The invention of the claims relates to a system and method using genetic algorithms for solving combinational multimodal optimization problems.

A listing of each independent claim, each dependent claim argued separately and each claim having means plus function language is provided below including exemplary references to page and line numbers of specification.

1. An automated computerized method for optimizing allocation of a set W of n tasks to m available resources for accomplishing such tasks using combinatorial multimodal optimization for finding multiple optimal ways of dividing said set *W* of *n* task values into *m* respectively groups associated with said resources, such that each of the groups satisfies a respective constraint condition, the method including execution of a computer program to automatically perform a series of machine operations comprising:

(a) receiving digital data signals representing plural tasks for assignment to available resources and, based thereon, defining an initial population of trial solutions assigning specific tasks to specific resources [p. 7, lns. 23-25; p. 8, ln. 12-24];

(b) calculating for each trial solution a fitness vector comprising *m* elements, each of which is indicative of whether the constraint condition of a corresponding respective one of the *m* groups has been satisfied by the trial solution [p. 7, lns. 26-27; p. 9, lns. 1-14];

(c) selecting a plurality of trial solutions for a next generation in dependence upon their respective fitness vectors [p. 8, lns. 1-3; p. 12, ln. 16 to p. 13, ln. 13];

(d) creating a new population of trial solutions including the selected earlier trial solutions [p. 8, lns. 4-7; p. 14, ln. 5 to p. 15, ln. 8];

(e) repeating steps (b) to (d) until the population of trial solutions stabilizes, the individual trial solutions of the stable population representing multiple optional ways of dividing the set *W* of tasks [p. 8, ln. 8]; and

(f) outputting at least one of said stabilized population as an optimized allocation of tasks to resources [p. 8, ln. 8].

14. An automated computerized method of distributing a plurality of tasks between a plurality of devices connected together to form a network, wherein each device has an associated constraint on the amount of tasks that it can perform per unit of time, the method including execution of a computer program to automatically perform a series of machine operations comprising:

(a) generating a plurality of trial solution allocations of tasks to devices to form an initial population of allocations [p. 7, lns. 23-25; p. 8, ln. 12-24];

(b) calculating for each trial solution a fitness vector comprising a plurality of elements each of which is indicative of whether the associated constraint of a corresponding respective one of the plurality of devices has been satisfied by the trial solution [p. 7, lns. 26-27; p. 9, lns. 1-14];

(c) selecting a plurality of allocations of tasks to devices for inclusion in a next generation of allocations in dependence upon their respective fitness vectors [p. 8, lns. 1-3; p. 12, ln. 16 to p. 13, ln. 13];

(d) creating the next generation of allocations of tasks to devices by including the allocations selected in step (c) together with new allocations, each of which is formed from a combination of two or more of the allocations selected in step (c) [p. 8, lns. 4-7; p. 14, ln. 5 to p. 15, ln. 8];

(e) repeating steps (b) to (d) until the population stabilizes [p. 8, ln. 8]; and

(f) outputting an allocation of the tasks among the devices according to one of the allocations included in the stabilized population [p. 8, ln. 8].

21. A tangible medium containing a computer program which, when executed effects a method for optimizing allocation of a set W of n tasks to m available resources for accomplishing such tasks using combinatorial multimodal optimization for finding multiple optimal ways of dividing said set W of n task values into m resource groups, such that each of the groups satisfies a respective constraint condition, the method including execution of a computer program to automatically perform a series of machine operations comprising:

(a) defining an initial population of trial solutions [p. 7, lns. 23-25; p. 8, ln. 12-24];

(b) calculating for each trial solution a fitness vector comprising *m* elements, each of which is indicative of whether the constraint condition of a corresponding respective one of the *m* groups has been satisfied by the trial solution [p. 7, lns. 26-27; p. 9, lns. 1-14];

(c) selecting a plurality of trial solutions for a next generation in dependence upon their respective fitness vectors [p. 8, lns. 1-3; p. 12, ln. 16 to p. 13, ln. 13];

(d) creating a new population of trial solutions including the selected earlier trial solutions [p. 8, lns. 4-7; p. 14, ln. 5 to p. 15, ln. 8];

(e) repeating steps (b) to (d) until the population of trial solutions stabilizes, the individual trial solution of the stable population representing multiple optional ways of dividing the set W of tasks [p. 8, ln. 8]; and

(f) outputting at least one of said stabilized population as an optimized allocation of tasks to resources [p. 8, ln. 8].

23. A system comprising a plurality of devices connected together to form a network, wherein each device has an associated constraint on the amount of tasks that it can perform per unit of time, the system including an allocation subsystem for allocating a plurality of tasks among the devices, the allocation subsystem comprising:

(a) means for generating a plurality of trial solution allocations to form an initial population of allocations [Figs. 1-3; p. 7, lns. 23-25; p. 8, ln. 12-24];

(b) means for calculating for each trial solution allocation a fitness vector comprising a plurality of elements each of which is indicative of whether the associated constraint of a corresponding respective one of the plurality of devices has been satisfied by the trial solution [Figs. 1-3; p. 7, lns. 26-27; p. 9, lns. 1-14];

(c) means for selecting a plurality of allocations for inclusion in a next generation of allocations in dependence upon their respective fitness vectors [Figs. 1-3; p. 8, lns. 1-3; p. 12, ln. 16 to p. 13, ln. 13];

(d) means for creating the next generation of allocations by including the allocations selected in step (c) together with new allocations each of which is formed

from a combination of two or more of the allocations selected in step (c) [Figs. 1-3; p. 8, lns. 4-7; p. 14, ln. 5 to p. 15, ln. 8];

(e) means for repeating steps (b) to (d) until the population stabilizes [Figs. 1-3; p. 8, ln. 8]; and

(f) means for outputting an allocation of the tasks among the devices according to one of the allocations included in the stabilized population [Figs. 1-3; p. 8, ln. 8].

24. A method of operating a multi-processor computer system to execute a computer program including a set of multiple separate tasks which must all be completed in order for the program execution to be complete, said method comprising:

distributing multiple of said set of program tasks between multiple computer program processor devices to efficiently accomplish all such distributed tasks wherein each computer program processor device has an associated constraint on the amount of tasks that it can perform per unit of time [p. 6, lns. 2-5], said distribution of tasks to said processor devices being accomplished by:

(a) receiving digital data signals representing a set of plural tasks for assignment to available processor devices and, based thereon, defining an initial population of trial solutions assigning specific tasks to specific processor devices p. 7, lns. 23-25; p. 8, ln. 12-24];

(b) calculating for each trial solution a fitness vector comprising a plurality of elements each of which is indicative of whether the constraint of a corresponding

respective one of the multiple computer program processor devices has been satisfied by the trial solution [p. 7, lns. 26-27; p. 9, lns. 1-14];

(c) selecting a plurality of trial solutions for a next generation in dependence upon their respective fitness vectors [p. 8, lns. 1-3; p. 12, ln. 16 to p. 13, ln. 13];

(d) creating a new population of trial solutions including the selected earlier trial solutions [p. 8, lns. 4-7; p. 14, ln. 5 to p. 15, ln. 8];

(e) repeating steps (b) to (d) until the population of trial solutions stabilizes, the individual trial solutions of the stable population representing multiple optional ways of dividing the input set of tasks [p. 8, ln. 8]; and

(f) outputting task assignments to said processor devices in conformance with at least one of said stabilized population as an optimized allocation of tasks to resources [p. 8, ln. 8].

25. A multi-processor computer system for executing a computer program including a set of multiple separate tasks which must all be completed in order for the program execution to be complete, said system comprising:

a plurality of computer program processors [Figs. 1-3; p. 18, lns. 1-20]; and

means networked with said multiple computer program processors for distributing multiple of said set of program tasks between said multiple computer program processor devices to efficiently accomplish all such distributed tasks wherein each computer program processor device has an associated constraint on the amount of tasks that it can perform per unit of time [Figs. 1-3; p. 18, lns. 1-20], said distribution of tasks to said processor devices being accomplished by:

(a) receiving digital data signals representing a set of plural tasks for assignment to available processors and, based thereon, defining an initial population trial solutions assigning specific tasks to specific processors p. 7, lns. 23-25; p. 8, ln. 12-24];

(b) calculating for each trial solution a fitness vector comprising a plurality of elements each of which is indicative of whether the constraint of a corresponding respective one of the multiple computer program processor devices has been satisfied by the trial solution [p. 7, lns. 26-27; p. 9, lns. 1-14];

(c) selecting a plurality of trial solutions for a next generation in dependence upon their respective fitness vectors [p. 8, lns. 1-3; p. 12, ln. 16 to p. 13, ln. 13];

(d) creating a new population of trial solutions including the selected earlier trial solutions [p. 8, lns. 4-7; p. 14, ln. 5 to p. 15, ln. 8];

(e) repeating steps (b) to (d) until the population of trial solutions stabilizes, the individual trial solutions of the stable population representing multiple optional ways of dividing the input set of tasks [p. 8, ln. 8], and

(f) outputting task assignments to said processors in conformance with at least one of said stabilized population as an optimized allocation of tasks to resources [p. 8, ln. 8].

(VI) GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

A. Whether claims 1, 3-6, 8-21, and 23-25 are rejected under 35 U.S.C. 102 (b) as being anticipated by Corne et al., PCT # WO 02/03716 A1 ("Corne").

B. Whether claim 7 would have been obvious under 35 U.S.C. 103(a) over Corne in view of Buczak et al. U.S. Patent No. 6,957,200 ("Buczak").

(VII) ARGUMENT

A. Claims 1, 3-6, 8-21, and 23-25 are not anticipated by Corne.

The Corne reference has been misapplied against independent claims 1, 14, 21 and 23-25 in that Corne discloses solving a different problem in a different way, as compared to the present claims. Indeed, at page 3 of the Office Action, the Examiner mistakenly asserts that in Corne "[t]he nodes in the network are resources." To the contrary, in Corne it is the links between the nodes which are the critical resources for the problem addressed, since each of the three objectives which Corne is seeking to optimize concerns the links rather than the nodes.

> The offline routing problem to be solved by embodiments of the present invention can be expressed as follows: to route multiple traffic requests r such that:
> a) no link is over-capacitated,
> b) communications costs associated with use of a link are minimised, and
> c) link utilizations are all below a specified, fixed target utlilisation,
> and therefore addresses 3 objectives of the offline routing problem.
> The bandwidth capacities of links in the network are of two types: a backbone type, having a capacity of 64 units (nodes 1 and 4 in Figure 2), and a local type having a capacity of 16 units (nodes 2, 3, 5-8).

See Corne at page 12, line 26 to page 13, line 1 (emphasis supplied). Thus, the stated objectives of Corne, i.e., objectives a) b) and c) identified above, all concern links.

Furthermore, Figure 2 of Corne illustrates the network whose routing is to be optimized as having 11 links (7 nodes) and further illustrates link costs 1-4. It is noteworthy that Corne defines its fitness vector as having three elements, and that each element of the fitness vector is an equation – namely, equations 1, 2 and 3 set out on

pages 13 and 14 of Corne – which can be evaluated for any given possible solution to arrive at three numbers which then form the elements for any given individual fitness vector for a given "solution."

> The three objective functions defined above thus provide individual components f1(x), f2(x), f3(x) (<u>in this case 3 because we have 3 objectives</u>) . . .

See Corne at page 14, lines 20-22 (emphasis supplied). Moreover, the equations are generally summations in respect of all of the links in the network so as to arrive at functions which are indicative of the global situation, rather than being indicative only of a local solution pertaining only to a single resource.

With respect to the rejected claims, the Examiner alleges that the claim term "groups" is not further defined in the claims and so that it is reasonable for the Examiner to read the "objectives" of Corne as equivalent to the groups required in the present claims. See Final Office Action at pages 3 and 11. Appellant disagrees with the Examiner's assertion. In fact, there are many limitations on the term "groups" throughout each of the present claims which make it improper for the Examiner to interpret the objectives of Corne as being in any way equivalent to the groups disclosed and claimed in the present application.[1]

First, the term "groups" is defined (and therefore limited) in multiple elements including the preamble of the relevant claims. For example, the preamble of claim 1 specifies a "method for optimizing the allocation of a set W of n tasks to m available

[1] Tellingly, the Examiner's cited portions of Corne do <u>not</u> make any mention of the term "groups." See Final Office at page 3 citing Corne at page 2, lines 4-6; page 3, line 30 to page 4, line 13; page 8, lines 26-32; page 9, lines 16-21, 23-30; page 12, line 26 to page 15, line 19; page 16, lines 5-23; page 24, line 12 to page 25, line 11; page 28, lines 14-19; and Figs. 3 and 4.

resources for accomplishing such tasks using combinatorial multimodal optimization for finding multiple optimal ways of dividing said set W of n tasks into m respectively groups associated with said resources" and the preamble of claim 21 specifies substantially the same method. This makes it clear that there are to be as many groups as there are resources (in particular m of each) and that each group is associated with a particular resource (hence the modifier "respectively"). This is a clear limitation on the meaning of the term "groups" which appears in these claims.

The claim term "groups" is further defined and hence limited at step (b) of claims 1 and 21 “calculating for each trial solution a fitness vector comprising m elements, each of which is indicative of whether the constraint condition of a corresponding respective one of the m groups has been satisfied by the trial solution." Use of the word "the" preceding the claim term "groups" makes clear that Applicant is referring back to the “groups” referred to in the preamble of the claims, as discussed above.[2] Furthermore, it is clear that there are m groups where m is the number of resources and it follows that there is a corresponding group for every resource and there are m of each, i.e., m resources and m groups with each group corresponding to a respective resource. Once an optimized solution has been found, each group of tasks will be assigned to its corresponding resource.

---

[2] In the Advisory Action, the Examiner alleges that because the preamble of claim 1 has the terms W, n and m italicized on one line and not on another line, that the terms must be different. The non-italicized terms were added by amendment on September 18, 2006 and are obviously a typographical error. Appellant respectfully requests that the correction of this typographical error be held in abeyance until this appeal has been decided at which time Appellant will correct the typographical error in accordance with the decision.

On the other hand, as noted above, while Corne discloses 11 links (and 7 nodes) how many objectives are there? There are 3 objective functions, but 3 does not equal 11 (nor 7). There is not a corresponding respective objective function associated with each resource – if there were there would have to be more than 3 objective functions since the number of resources is 11 (if you take the resources to equal the links, or 7 if you take the resources to be nodes as the Examiner has suggested). Since each "group" of the present claims has a corresponding respective resource and since each objective function of Corne does not have an associated resource it stands to reason that the objectives of Corne do not correspond to the claimed groups of Applicant's invention.

Furthermore, even if the objectives of Corne did correspond to the groups required in the present claims, the present claims further require that the fitness vector has m elements – however, in Corne the fitness vector has only 3 elements whereas m (which is equal to the number of resources) equals 11 (if the resources are links, or 7 if the resources ore nodes) so Corne also does not satisfy the requirement in the present claims that the fitness vector has m elements.

The above arguments have been specifically addressed to independent claims 1 and 21 which expressly recite the claim term "groups." However, the same arguments are applicable to independent claims 14 and 23-25 which require ". . . elements each of which is indicative of whether the associated constraint of a corresponding respective one of the [plurality of devices **or** multiple computer processor devices] has been satisfied by the trial solution." This requires that there are (at least) as many elements as there are "resources" where the "resource" in question depends upon the particular

claim (e.g., in claims 14 and 23 its "devices" in claims 24 and 25 its "processor devices"). Again, in Corne there are only three objectives and each fitness vector correspondingly has only 3 elements and so Corne does not have "elements each of which is indicative of whether the constraint of a corresponding one of the multiple devices/processor devices . . . " rather each element in Corne is indicative of the global extent to which a particular global objective function has been met. Thus the same arguments apply for independent claims 14 and 23-25 as for independent claims 1 and 21.

For all of the above reasons, it is respectfully submitted that Corne does not anticipate the present claims. Appellant's claimed solution provides a useful technique for approaching a certain type of problem which is not taught or even suggested in Corne. Accordingly, independent claims 1, 14, 21 and 23-25 and their respective dependents claims patentably define over Corne.

B. <u>Claim 7 would not have been obvious over Corne in view of Buczak</u>.

The Examiner has cited Buczak merely for teaching that a non-reserved proportion of the new population is generated using a Roulette wheel selection method. Accordingly, it should be clear that Buczak does not overcome the deficiencies noted above with respect to Corne. Therefore, all of the present claims are believed to patentably define over the cited art taken singly or in combination.

CONCLUSION

In conclusion it is believed that the application is in clear condition for allowance; therefore, early reversal of the Final Rejection and passage of the subject application to issue are earnestly solicited.

Respectfully submitted,

NIXON & VANDERHYE P.C.

By: ____________________
Chris Comuntzis
Reg. No. 31,097

CC:lmr
901 North Glebe Road, 11th Floor
Arlington, VA 22203-1808
Telephone: (703) 816-4000
Facsimile: (703) 816-4100

(VIII) CLAIMS APPENDIX

1. An automated computerized method for optimizing allocation of a set W of n tasks to m available resources for accomplishing such tasks using combinatorial multimodal optimization for finding multiple optimal ways of dividing said set *W* of *n* task values into *m* respectively groups associated with said resources, such that each of the groups satisfies a respective constraint condition, the method including execution of a computer program to automatically perform a series of machine operations comprising:

(a) receiving digital data signals representing plural tasks for assignment to available resources and, based thereon, defining an initial population of trial solutions assigning specific tasks to specific resources;

(b) calculating for each trial solution a fitness vector comprising *m* elements, each of which is indicative of whether the constraint condition of a corresponding respective one of the *m* groups has been satisfied by the trial solution;

(c) selecting a plurality of trial solutions for a next generation in dependence upon their respective fitness vectors;

(d) creating a new population of trial solutions including the selected earlier trial solutions;

(e) repeating steps (b) to (d) until the population of trial solutions stabilizes, the individual trial solutions of the stable population representing multiple optional ways of dividing the set *W* of tasks; and

(f) outputting at least one of said stabilized population as an optimized allocation of tasks to resources.

Claim 2 (Canceled).

3. A method as in claim 1 in which the fitness vector comprises *m* bits, each bit being indicative of whether the constraint condition of a corresponding one of the *m* groups has been satisfied.

4. A method as in claim 1 including calculating a fitness value for each individual trial solution.

5. A method as claimed in claim 3 including calculating a fitness value for each individual trial solution in which the fitness value comprises the sum of the bits in the fitness vector.

6. A method as in claim 1 including reserving a proportion of the new population for individual trial solutions selected at step (c).

7. A method as 6 in which a non-reserved proportion of the new population is generated using a Roulette wheel selection method.

8. A method as in claim 1 in which step (c) comprises selecting non-dominated individual trial solutions using the criteria of Pareto optimality.

9. A method as in claim 4 in which step (c) comprises selecting non-dominated individual trial solutions using the criteria of Pareto optimality including ranking non-dominated individual trial solutions by fitness value, and selecting from the ranked list.

10. A method as in claim 9 in which only non-dominated individual trial solutions with greatest fitness value may be selected at step (c).

11. A method as in claim 4 in which step (c) comprises selecting individual trial solutions in dependence upon both their respective fitness vectors and their respective fitness values.

12. A method as in claim 1 in which crossover and mutation are applied at step (d) to at least some individual trial solutions in the new population.

13. A method as in claim 1 in which step (c) comprises selecting no more than one individual trial solution for each unique fitness vector.

14. An automated computerized method of distributing a plurality of tasks between a plurality of devices connected together to form a network, wherein each device has an associated constraint on the amount of tasks that it can perform per unit of time, the method including execution of a computer program to automatically perform a series of machine operations comprising:

(a) generating a plurality of trial solution allocations of tasks to devices to form an initial population of allocations;

(b) calculating for each trial solution a fitness vector comprising a plurality of elements each of which is indicative of whether the associated constraint of a corresponding respective one of the plurality of devices has been satisfied by the trial solution;

(c) selecting a plurality of allocations of tasks to devices for inclusion in a next generation of allocations in dependence upon their respective fitness vectors;

(d) creating the next generation of allocations of tasks to devices by including the allocations selected in step (c) together with new allocations, each of which is formed from a combination of two or more of the allocations selected in step (c);

(e) repeating steps (b) to (d) until the population stabilizes; and

(f) outputting an allocation of the tasks among the devices according to one of the allocations included in the stabilized population.

15. A method as in claim 14 wherein the devices are processors within a multi-processor computer system.

16. A method as in claim 14 wherein the devices are computers within a computer network.

17. A method as in claim 14 wherein the devices are routers and the tasks are estimated volumes of traffic to be routed through the routers within a data network, and wherein the allocations are used to form a routing strategy.

18. A method as in claim 14 in which step (c) comprises selecting non-dominated allocations using the criteria of Pareto optimality of the associated fitness vectors.

19. A method as in claim 1 in which new allocations are formed in step (d) by performing crossover operations in respect of groups of two or more of the allocations selected in step (c).

20. A method as in claim 14 in which mutation operations are applied to one or more of the new allocations formed in step (d) according to a predetermined probability of each new allocation being mutated.

21. A tangible medium containing a computer program which, when executed effects a method for optimizing allocation of a set W of n tasks to m available resources for accomplishing such tasks using combinatorial multimodal optimization for finding multiple optimal ways of dividing said set W of n task values into m resource groups, such that each of the groups satisfies a respective constraint condition, the method including execution of a computer program to automatically perform a series of machine operations comprising:

(a) defining an initial population of trial solutions;

(b) calculating for each trial solution a fitness vector comprising *m* elements, each of which is indicative of whether the constraint condition of a corresponding respective one of the *m* groups has been satisfied by the trial solution;

(c) selecting a plurality of trial solutions for a next generation in dependence upon their respective fitness vectors;

(d) creating a new population of trial solutions including the selected earlier trial solutions;

(e) repeating steps (b) to (d) until the population of trial solutions stabilizes, the individual trial solution of the stable population representing multiple optional ways of dividing the set W of tasks; and

(f) outputting at least one of said stabilized population as an optimized allocation of tasks to resources.

Claim 22 (Canceled).

23. A system comprising a plurality of devices connected together to form a network, wherein each device has an associated constraint on the amount of tasks that it can perform per unit of time, the system including an allocation subsystem for allocating a plurality of tasks among the devices, the allocation subsystem comprising:

(a) means for generating a plurality of trial solution allocations to form an initial population of allocations;

(b) means for calculating for each trial solution allocation a fitness vector comprising a plurality of elements each of which is indicative of whether the associated constraint of a corresponding respective one of the plurality of devices has been satisfied by the trial solution;

(c) means for selecting a plurality of allocations for inclusion in a next generation of allocations in dependence upon their respective fitness vectors;

(d) means for creating the next generation of allocations by including the allocations selected in step (c) together with new allocations each of which is formed from a combination of two or more of the allocations selected in step (c);

(e) means for repeating steps (b) to (d) until the population stabilizes; and

(f) means for outputting an allocation of the tasks among the devices according to one of the allocations included in the stabilized population.

24. A method of operating a multi-processor computer system to execute a computer program including a set of multiple separate tasks which must all be completed in order for the program execution to be complete, said method comprising:

distributing multiple of said set of program tasks between multiple computer program processor devices to efficiently accomplish all such distributed tasks wherein each computer program processor device has an associated constraint on the amount of tasks that it can perform per unit of time, said distribution of tasks to said processor devices being accomplished by:

(a) receiving digital data signals representing a set of plural tasks for assignment to available processor devices and, based thereon, defining an initial population of trial solutions assigning specific tasks to specific processor devices;

(b) calculating for each trial solution a fitness vector comprising a plurality of elements each of which is indicative of whether the constraint of a corresponding respective one of the multiple computer program processor devices has been satisfied by the trial solution;

(c) selecting a plurality of trial solutions for a next generation in dependence upon their respective fitness vectors;

(d) creating a new population of trial solutions including the selected earlier trial solutions;

(e) repeating steps (b) to (d) until the population of trial solutions stabilizes, the individual trial solutions of the stable population representing multiple optional ways of dividing the input set of tasks; and

(f) outputting task assignments to said processor devices in conformance with at least one of said stabilized population as an optimized allocation of tasks to resources.

25. A multi-processor computer system for executing a computer program including a set of multiple separate tasks which must all be completed in order for the program execution to be complete, said system comprising:

a plurality of computer program processors; and

means networked with said multiple computer program processors for distributing multiple of said set of program tasks between said multiple computer program processor devices to efficiently accomplish all such distributed tasks wherein each computer program processor device has an associated constraint on the amount of tasks that it can perform per unit of time, said distribution of tasks to said processor devices being accomplished by:

(a) receiving digital data signals representing a set of plural tasks for assignment to available processors and, based thereon, defining an initial population trial solutions assigning specific tasks to specific processors;

(b) calculating for each trial solution a fitness vector comprising a plurality of elements each of which is indicative of whether the constraint of a corresponding

respective one of the multiple computer program processor devices has been satisfied by the trial solution;

(c) selecting a plurality of trial solutions for a next generation in dependence upon their respective fitness vectors;

(d) creating a new population of trial solutions including the selected earlier trial solutions;

(e) repeating steps (b) to (d) until the population of trial solutions stabilizes, the individual trial solutions of the stable population representing multiple optional ways of dividing the input set of tasks, and

(f) outputting task assignments to said processors in conformance with at least one of said stabilized population as an optimized allocation of tasks to resources.

(IX) EVIDENCE APPENDIX

None.

(X) RELATED PROCEEDINGS APPENDIX

None.